Empirical best prediction for small area estimation using nonparametric maximum likelihood

M.F. Marino* M.G. Ranalli † N. Salvati ‡ M. Alfò §

April 10, 2017

arXiv:1704.02220v1 [stat.ME] 7 Apr 2017

Abstract

In small area estimation, generalized linear mixed models represent a useful tool for deriving best prediction of counts or proportions. For non-Gaussian responses, computing the Empirical Best Prediction and the corresponding analytic approximation to its Mean Squared Error requires the solution of (possibly) multiple integrals that, in general, do not admit closed form. Monte Carlo methods and parametric bootstrap are common choices, even though their computational burden represents a non trivial issue. We propose to estimate model parameters within a NonParametric Maximum Likelihood framework. In this context, the distribution of the area-specific random parameters is left unspecified and is approximated by a (discrete) nonparametric distribution. Given the discrete nature of the mixing distribution, we can avoid integral approximations and considerably reduce the computational effort in the case of non-Gaussian responses. Within this framework, we derive the Empirical Best Prediction for a (possibly non-linear) mixed effect and the analytic approximation of the corresponding MSE. The proposed approach is presented for a general response that belongs to the Exponential Family. Then, we focus on the relevant case of binary data. The method is tested via a large scale simulation study and applied to unit-level data from the 2012 Italian Labor Force Survey.

Key Words: NPML; Exponential Family; Finite Mixture; Unit-level model; EM algorithm.

1 Introduction

Small area estimation has received considerable attention in the past decades both in terms of theoretical developments and applications in the Official Statistics context. Inference in this framework refers to the issue of obtaining estimates for various characteristics of interest referring to sub-populations with limited sample size. Direct, usually design-based, estimates have unduly large sampling errors and indirect, model based, predictors are typically used, even if possibly biased, because of their much smaller error. An updated appraisal is given by Rao and Molina (2015).

Linear mixed models represent a typical tool when dealing with small area estimation for deriving the best prediction of totals and means for a continuous response. Area-specific random

*Dep. of Political Science, Università degli Studi di Perugia, Italy. mariafrancesca.marino@unipg.it
†Dep. of Political Science, Università degli Studi di Perugia, Italy. giovanna.ranalli@unipg.it
‡Dep. of Economics and Management, Università di Pisa, Italy. nicola.salvati@unipi.it
§Dep. of Statistical Science, Sapienza Università di Roma, Italy. marco.alfo@uniroma1.it

parameters are usually included in the model to account for sources of unobserved heterogeneity that are not captured by the covariates and, additionally, to allow for correlation between units within the same area. For Gaussian data, Prasad and Rao (1990) developed an Empirical Best Linear Unbiased Predictor (EBLUP) to estimate small area characteristics. Similarly, Jiang and Lahiri (2001) developed an Empirical Best Prediction (EBP) method for estimating small area proportions under a mixed logistic model. Jiang (2003) further extended this approach to general responses belonging to the Exponential Family within a generalized linear mixed model (GLMM) framework. More recently, Boubeta et al. (2016) derived the EBP under an area-level mixed Poisson model for small area counts, while Hobza and Morales (2016) investigated the EBP and the corresponding mean squared error (MSE) estimators under a unit-level mixed logistic model. Generally, computing the EBP and evaluating its precision for non-Gaussian data is not an easy task. As highlighted by Hobza and Morales (2016) and Boubeta et al. (2016), the main drawback is related to the computational burden, which makes the method difficult to be implemented in most real-world applications. For this reason, ad-hoc alternatives, mainly based on plug-in predictors and Taylor linearization, have been proposed and are currently largely applied. See e.g. Saei and Chambers (2003), González-Manteiga et al. (2007), and Molina et al. (2007).

In all the aforementioned approaches, the random parameter distribution is assumed to be Gaussian. This assumption is not essential to compute the EBP as, for this purpose, we only need the distribution of the random parameters to be specified up to a vector of unknown parameters (Jiang and Lahiri, 2001). Here, we propose to make inference on GLMM parameters exploiting a NonParametric Maximum Likelihood (NPML) approach (Aitkin, 1996, 1999). In this framework, the distribution of the area specific random parameters is left completely unspecified and is estimate directly from the observed data. Such a method is known to lead to a discrete distribution defined over a pre-specified (relatively small) number of locations (Lindsay, 1983a). The resulting likelihood resembles that of a finite mixture of regression models for responses belonging to the Exponential Family and, thus, this approach leads to a finite mixture of generalized linear models. Starting from the unspecified mixing distribution for the random parameters, we introduce a nonparametric version of the EBP for small area quantities of interest and a second order approximation of the corresponding MSE. The history of NPML is quite long and dates back, at least, to Kiefer and Wolfowitz (1956). With increasing computational resources and due to the results of Simar (1976), Laird (1978), Böhning (1982), and Lindsay (1983a,b), the NPML approach has been given an increasing interest. In particular, it has been used to account for overdispersion (e.g. Hinde and Wood, 1987; Aitkin, 1996), clustering (e.g. Wedel et al., 1993) and longitudinal data (e.g. Aitkin, 1999; Aitkin and Alfò, 2003; Alfò and

Aitkin, 2006), also in the context of quantile regression (Alfò et al., 2017). Nonparametric methods have also been introduced in small area estimation to handle departures from linearity in the relationship with one or more covariates using penalized splines (Opsomer et al., 2008).

The proposed approach offers a number of advantages. First, it allows us to avoid unverifiable assumptions on the random parameter distribution. Second, since mixture parameters are directly estimated from the data and are completely free to vary over the corresponding support, extreme and/or asymmetric departures from the homogeneous model can be easily accommodated. Last and more important, the discrete nature of the mixing distribution allows us to avoid numerical approximations, as those based on (adaptive) Gaussian quadrature or Laplace approximations (see e.g. Pinheiro and Bates, 1995), or simulation based approaches, as those using Monte Carlo approximations (see e.g. McCulloch, 1997) for deriving parameter estimates in a maximum likelihood (ML) framework, for computing EBPs of nonlinear functions of mixed effects, and evaluating the corresponding MSE. As a result, we may considerably reduce the corresponding computational effort required when dealing with non-Gaussian responses.

The computational gains obtained using the NPML approach in the small area framework are particularly evident when addressing the issue of measuring the accuracy of EBPs. Estimation of the MSE for the EBP in the case of non-Gaussian responses is not an easy task. Jiang and Lahiri (2001) studied the analytic approximation of the MSE for the EBP of binary data and provided a second-order, bias corrected, estimator. Jiang (2003) generalized such an approximation for general responses in the Exponential Family. Recently, Boubeta et al. (2016) and Hobza and Morales (2016) used this approach in the case of area-level mixed Poisson models and unit-level mixed logistic models, respectively. In this paper, we adapt the calculations in Jiang and Lahiri (2001) and Jiang (2003) to obtain an analytic approximation of the MSE for the nonparametric EBP (np-EBP) in the case of a finite mixture of generalized linear models.

We present the proposed approach for a general response with density in the Exponential Family and, later, we focus on the relevant case of binary data. We carried out a series of simulation studies to investigate the performance of the proposed np-EBP, focusing on its prediction accuracy and coverage. We compared our proposal to the EBP approach by Jiang and Lahiri (2001) in terms of both prediction accuracy and computational burden and to the plug-in estimator by e.g. Saei and Chambers (2003) and González-Manteiga et al. (2007). We also applied the proposed np-EBP approach to unit-level data from the Italian Labor Force Survey (ILFS), with the aim of estimating unemployment incidence in Italy.

The paper is organized as follows. Section 2 introduces notation and a brief review of the EBP and its MSE approximation as proposed by Jiang (2003). In Section 3 we describe the NPML approach and the corresponding estimation of model parameters. The np-EBP and the

corresponding MSE approximation we propose are described in Section 4. Section 5 focuses on the particular case of binary responses. Section 6 reports the results of the simulation studies, while Section 7 entails the application of the proposed approach to the ILFS data. Finally, Section 8 summarizes our findings and provides guidelines for future research.

2 Empirical best prediction in GLMMs

Let U denote a finite population of size N which can be partitioned into m non-overlapping small areas or domains, with U_i denoting small area i. The small area population sizes N_i, $i = 1, \ldots, m$, are assumed known. Let n denote the overall sample size drawn from the above population. The set s_i contains the n_i population indices of the sample units in small area i, with $n = \sum_{i=1}^{m} n_i$.

Let Y_{ij} and $\boldsymbol{x}_{ij} = (x_{ij1}, \ldots, x_{ijp})'$ denote a response variable and a set of p covariates recorded for the j-th unit belonging to the i-th small area, $i = 1, \ldots, m$, $j = 1, \ldots, N_i$. Furthermore, let $\boldsymbol{\alpha}_i$ be a zero mean, area-specific, q-dimensional, random parameter vector, with $q \leq p$, distributed according to the density $f_\alpha(\cdot)$. Let us assume, this latter is indexed by the unknown covariance matrix $\boldsymbol{\Sigma}$. Last, let $\boldsymbol{w}_{ij}$ denote a q-dimensional subset of covariates out of the p in $\boldsymbol{x}_{ij}$ that are associated to the random parameters $\boldsymbol{\alpha}_i$. In the following, we will assume that the values of $\boldsymbol{x}_{ij}$, and therefore of $\boldsymbol{w}_{ij}$, are known for all units in the population. This can be quite a restrictive assumption in some real-world applications, since it implies the availability of individual census records. When the auxiliary variables are categorical and/or take a finite number of values this assumption can be relaxed. We will discuss this issue in more detail when dealing with the application to the ILFS data.

2.1 The model

According to a local independence assumption, we assume that, conditional on the area-specific random parameters $\boldsymbol{\alpha}_i$, responses $Y_{ij}, i = 1, \ldots, m, j = 1, \ldots, N_i$, are independent and follow an Exponential Family distribution with density

$$f_{y|\alpha}(y_{ij} \mid \boldsymbol{\alpha}_i) = \exp\left\{\frac{y_{ij}\theta_{ij} - b(\theta_{ij})}{a(\phi)} + c(y_{ij}, \phi)\right\}.$$

In the previous expression, ϕ is a dispersion parameter, $a(\cdot), b(\cdot)$ and $c(\cdot)$ are known functions and θ_{ij} is the canonical parameter specific to the chosen member of the family. Let us denote by $\boldsymbol{\beta}$ a p-dimensional vector of fixed regression coefficients and let us assume that θ_{ij} is associated to the following mixed model (Laird and Ware, 1982):

$$\eta_{ij} = \boldsymbol{x}'_{ij}\boldsymbol{\beta} + \boldsymbol{w}'_{ij}\boldsymbol{\alpha}_i,$$

through the following link function: $g[\theta_{ij}] = \eta_{ij}$. When a canonical link is employed, $\theta_{ij} = \eta_{ij}$. Therefore, the joint distribution of the response vector $\boldsymbol{y}_i = (y_{i1}, \ldots, y_{iN_i})'$ associated to the i-th small area, conditional on the area-specific random parameter vector $\boldsymbol{\alpha}_i$, is given by

$$f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_i; \boldsymbol{X}_i) = \exp\left\{ \sum_{j=1}^{N_i} \frac{y_{ij}\,\theta_{ij} - b(\eta_{ij})}{a(\phi)} + c(y_{ij}, \phi) \right\},$$

where $\boldsymbol{X}_i$ denotes the matrix of all covariates associated to the units in area i. The marginal distribution of area-specific responses $\boldsymbol{y}_i$ is obtained by integrating out the corresponding random effect, $\boldsymbol{\alpha}_i$, that is

$$f_y(\boldsymbol{y}_i) = \int_{\mathbb{R}^q} f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_i; \boldsymbol{X}_i) f_\alpha(\boldsymbol{\alpha}_i) d\boldsymbol{\alpha}_i.$$

Typically, a parametric specification for $f_\alpha(\boldsymbol{\alpha}_i)$ is adopted, and a zero mean multivariate Gaussian distribution is a common choice. Also, it is worth noticing that an implicit assumption regarding the random parameter distribution is that of (strong) exogeneity of observed covariates $\boldsymbol{x}_{ij}$, that is $f_\alpha(\boldsymbol{\alpha}_i \mid \boldsymbol{X}_i) = f_\alpha(\boldsymbol{\alpha}_i)$ or, more specifically, $E(\boldsymbol{\alpha}_i \mid \boldsymbol{X}_i) = E(\boldsymbol{\alpha}_i) = \boldsymbol{0}$. While this assumption may lead to inconsistent estimators if it does not hold, a simple correction which may be referred to as QP decomposition (Krishnakumar, 2006) can be adopted. This slightly modifies the linear predictor above and produces area-specific random parameter that are (linearly) free of $\boldsymbol{X}_i$, see Neuhaus and McCulloch (2006). Therefore, we will proceed by assuming that, if needed, such a correction is applied and that $f_\alpha(\boldsymbol{\alpha}_i \mid \boldsymbol{X}_i) = f_\alpha(\boldsymbol{\alpha}_i)$.

2.2 The EBP and the MSE approximation

Let us assume, we are interested in using sample data on responses $Y_{ij}, i = 1, \ldots, m, j \in s_i$, and population data on covariates $\boldsymbol{x}_{ij}$ to predict a (possibly non-linear) function of fixed and random parameters, say $\zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma})$, where $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_1', \ldots, \boldsymbol{\alpha}_m')'$. According to Jiang (2003), the best predictor (BP) of ζ in terms of minimum MSE is given by

$$\tilde{\zeta}^{BP}(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}) = E_{\alpha|y}\left[\zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}) \mid \boldsymbol{y}\right] = \int_{\mathbb{R}^{m\times q}} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}) f_{\alpha|y}(\boldsymbol{\alpha} \mid \boldsymbol{y}) d\boldsymbol{\alpha}, \tag{1}$$

where

$$f_{\alpha|y}(\boldsymbol{\alpha} \mid \boldsymbol{y}) = \frac{\prod_{i=1}^m f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_i; \boldsymbol{X}_i) f_\alpha(\boldsymbol{\alpha}_i)}{\prod_{i=1}^m \int_{\mathbb{R}^q} f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_i; \boldsymbol{X}_i) f_\alpha(\boldsymbol{\alpha}_i) d\boldsymbol{\alpha}_i}.$$

In practice, model parameters are unknown and, therefore, need to be estimated.

Let $\boldsymbol{\Phi} = (\boldsymbol{\beta}, \phi, \boldsymbol{\Sigma})$ be the vector of model parameters; estimation can be accomplished by maximizing the following observed data likelihood function:

$$L(\boldsymbol{\Phi}) = \prod_{i=1}^m f_y(\boldsymbol{y}_i; \boldsymbol{X}_i) = \prod_{i=1}^m \int_{\mathbb{R}^q} f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_i; \boldsymbol{X}_i) f_\alpha(\boldsymbol{\alpha}_i) d\boldsymbol{\alpha}_i. \tag{2}$$

It is worth noticing that, in order to maximize equation (2), we first need to evaluate an integral defined over the support of the area-specific random parameters and this can be directly done in few cases, for instance when $f_{y|\alpha}(\cdot \mid \cdot)$ and $f_\alpha(\cdot)$ are conjugate. In all other cases, numerical approximations (e.g. Gaussian quadrature techniques) or simulation based methods (e.g. Monte Carlo integration) need to be used, but the computational complexity of such methods represent a non-trivial drawback. To overcome the issue, Jiang (1998) suggested to derive estimates by exploiting the method of moments. Penalized quasi-likelihood approach (e.g., Breslow and Clayton, 1993) represents a further alternative that is less computationally demanding. However, it is well-known that it may provide inconsistent estimates. In Section 3, we describe a different approach based on the use of area-specific random effects with an unspecified distribution. Once parameter estimates have been obtained, we can use them to compute the empirical best predictor (EBP) of ζ, that is $\hat{\zeta}^{EBP} = \tilde{\zeta}^{BP}(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Sigma}})$.

To evaluate the quality of predictions obtained via the EBP method, we focus on the estimation of the corresponding MSE. Following similar arguments to those detailed in Jiang (2003), it is easy to verify the following decomposition:

$$\text{MSE}(\hat{\zeta}^{EBP}) = E[(\hat{\zeta}^{EBP} - \zeta)^2] = E[(\hat{\zeta}^{EBP} - \tilde{\zeta}^{BP})^2] + E[(\tilde{\zeta}^{BP} - \zeta)^2]. \tag{3}$$

The second term in the above expression corresponds to the MSE of the best prediction $\tilde{\zeta}^{EB}$ and can be further decomposed as follows:

$$\begin{aligned}
\text{MSE}(\tilde{\zeta}^{BP}) &= E[(\tilde{\zeta}^{BP} - \zeta)^2] = E_\alpha[(\tilde{\zeta}^{BP})^2] + E_\alpha[(\zeta)^2] - 2E_\alpha[\tilde{\zeta}^{BP}\zeta] \\
&= E_y\{E_{\alpha|y}[(\tilde{\zeta}^{BP})^2]\} + E_\alpha[(\zeta)^2] - 2E_y\{E_{\alpha|y}[(\tilde{\zeta}^{BP}\,\zeta)]\} \\
&= E_y[(\tilde{\zeta}^{BP})^2] + E_\alpha[(\zeta)^2] - 2E_y\{\tilde{\zeta}^{BP}E_{\alpha|y}[(\zeta)]\} \\
&= E_y[(\tilde{\zeta}^{BP})^2] + E_\alpha[(\zeta)^2] - 2E_y[(\tilde{\zeta}^{BP})^2] \\
&= E_\alpha[(\zeta)^2] - E_y[(\tilde{\zeta}^{BP})^2] = d(\boldsymbol{\Phi}).
\end{aligned}$$

In particular, to derive $d(\boldsymbol{\Phi})$, we need to evaluate the following integrals:

$$E_\alpha[(\zeta)^2] = \int_{\mathbb{R}^{m\times q}} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma})^2 \, f_\alpha(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}, \tag{4}$$

$$E_y[(\tilde{\zeta}^{BP})^2] = E_y\left[\left(\int_{\mathbb{R}^{m\times q}} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}) f_{\alpha|y}(\boldsymbol{\alpha} \mid \boldsymbol{y})\, d\boldsymbol{\alpha}\right)^2\right], \tag{5}$$

where $f_\alpha(\boldsymbol{\alpha}) = \prod_{i=1}^m f_\alpha(\boldsymbol{\alpha}_i)$ denotes the joint density of all random parameters $\boldsymbol{\alpha}_i, i = 1, \ldots, m$. As it has been remarked before, the corresponding computational burden could represent a non-trivial issue.

Jiang (2003) also showed that, under mild regularity conditions, the first term on the right hand side of equation (3) can be approximated as follows

$$E[(\hat{\zeta}^{EBP} - \tilde{\zeta}^{BP})^2] = \frac{1}{m} E_y \left[\left(\frac{\partial \tilde{\zeta}^{BP}}{\partial \boldsymbol{\Phi}} \right)' m V(\hat{\boldsymbol{\Phi}}) \left(\frac{\partial \tilde{\zeta}^{BP}}{\partial \boldsymbol{\Phi}} \right) \right] + o_p(1/m)$$
$$= \frac{1}{m} e(\boldsymbol{\Phi}) + o_p\left(\frac{1}{m}\right), \tag{6}$$

where $\partial \tilde{\zeta}_{ij}^{BP} / \partial \boldsymbol{\Phi}$ denotes the first derivative vector of the BP reported in equation (1), while $V(\hat{\boldsymbol{\Phi}})$ denotes the covariance matrix of model parameter estimates. An estimator of $\text{MSE}(\hat{\zeta}^{EBP})$ can be obtained by replacing $\boldsymbol{\Phi}$ in equation(3) by a consistent estimator, that is

$$\widehat{\text{MSE}}(\hat{\zeta}^{EBP}) = d(\hat{\boldsymbol{\Phi}}) + \frac{1}{m} e(\hat{\boldsymbol{\Phi}}).$$

However, as discussed by Jiang (2003), while we may simply replace parameter estimates into $e(\boldsymbol{\Phi})$ as we get an error of order $o_p(1/m)$, the same result is not valid for $d(\boldsymbol{\Phi})$. In fact, under mild regularity conditions, the authors showed that

$$E[d(\hat{\boldsymbol{\Phi}})] = d(\boldsymbol{\Phi}) + \frac{1}{m} b(\boldsymbol{\Phi}) + o_p\left(\frac{1}{m}\right),$$

where

$$b(\boldsymbol{\Phi}) = \left(\frac{\partial d}{\partial \boldsymbol{\Phi}} \right)' m E(\hat{\boldsymbol{\Phi}} - \boldsymbol{\Phi}) + \frac{1}{2} E\left[\sqrt{m}(\hat{\boldsymbol{\Phi}} - \boldsymbol{\Phi})' \left(\frac{\partial^2 d}{\partial \boldsymbol{\Phi} \boldsymbol{\Phi}'} \right) \sqrt{m}(\hat{\boldsymbol{\Phi}} - \boldsymbol{\Phi}) \right], \tag{7}$$

represents a bias correction term, which can be used to modify $d(\hat{\boldsymbol{\Phi}})$ to get an unbiased estimator of $d(\boldsymbol{\Phi})$. As it is clear for the previous formulas, computing the BP, the EBP and the corresponding MSE requires the solution of (multiple) integrals that, generally, do no admit a closed form expression when dealing with non-Gaussian responses. As stated before, Monte Carlo approximations or numerical integration techniques are required and this makes the computation extremely time-consuming, especially when the derivation of the MSE in equation (3) is concerned. For these reasons, non-optimal estimators are often used for the prediction of small area quantities. Naive predictions based on Penalized Quasi Likelihood (PQL) estimates represent a common strategy: model parameter estimates are computed via a PQL approach and these are simply substituted in the ζ function one is interested in for prediction (see e.g. González-Manteiga et al., 2007). Similarly, a bootstrap approach is frequently employed for MSE estimation, especially when dealing with a limited number of small areas. However, when m is large, as in the ILFS dataset we will analyze in Section 7, neither the analytic approach by Jiang (2003) nor the bootstrap approach (Boubeta et al., 2016; Hobza and Morales, 2016) represent viable strategies due to computational issues. In the following sections, we will describe an alternative approach which allows us to considerably reduce the computational burden for deriving the above quantities, while preserving optimality of the results.

3 NonParametric Maximum Likelihood estimation

As highlighted before, although ML estimators derived by maximizing the observed data likelihood in equation (2) are optimal, their computation remains a non trivial matter when dealing with non-Gaussian responses. This is due to the presence of q-dimensional, often intractable, integrals which need to be numerically approximated to get estimates. As an alternative, we propose to follow the approach introduced by Aitkin (1996, 1999), and leave the distribution of $\boldsymbol{\alpha}_i$ completely unspecified. In particular, the area-specific random parameters are treated as nuisance parameters, the corresponding distribution is left unspecified and is estimated via a NonParametric Maximum Likelihood (NPML) approach. Different contributions to this theory can be found in the literature (Simar, 1976; Laird, 1978; Böhning, 1982; Lindsay, 1983a,b). In particular, Lindsay (1983a) showed that ML estimation of a mixing distribution reduces to a standard problem of maximizing a concave function over a convex set. As long as the likelihood function is bounded, it can be shown that $L(\boldsymbol{\Phi})$ can be maximized by a discrete distribution defined on at most $G \leq K$ support points, where K is the number of distinct area profiles in the sample. Let us assume that the mixing distribution puts masses $\pi_g > 0$ on locations $\boldsymbol{\alpha}_g = (\alpha_{g1}, \ldots, \alpha_{gq})', g = 1, \ldots, G$, where the constraint $\sum_{g=1}^{G} \pi_g = 1$ holds. Parameter estimates can be obtained by maximizing the following expression:

$$L(\boldsymbol{\Phi}) = \prod_{i=1}^{m} \sum_{g=1}^{G} f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_g; \boldsymbol{X}_i)\pi_g. \tag{8}$$

Here, $f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_g; \boldsymbol{X}_i)$ denotes the product of Exponential Family densities $f_{y|\alpha}(y_{ij} \mid \boldsymbol{\alpha}_g; \boldsymbol{x}_{ij})$ with canonical parameter θ_{ijg} which is defined according to the following mixed model:

$$g[\theta_{ijg}] = \eta_{ijg} = \boldsymbol{x}'_{ij}\boldsymbol{\beta} + \boldsymbol{w}'_{ij}\boldsymbol{\alpha}_g,$$

while $\boldsymbol{\Phi}$ denotes the global set of model parameters, that is $\boldsymbol{\Phi} = \{\boldsymbol{\beta}, \phi, \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_G, \pi_1, \ldots, \pi_G\}$. As it is clear, expression (8) resembles the likelihood of a finite mixture of distributions with weights $\pi_g = \Pr(\boldsymbol{\alpha}_i = \boldsymbol{\alpha}_g)$. That is, $\boldsymbol{\alpha}_i \sim \sum_{g=1}^{G} \pi_g \delta(\boldsymbol{\alpha}_g)$, where $\delta(a)$ is a one-point distribution putting a unit mass at a.

3.1 Parameter estimation

According to modeling assumptions, locations $\boldsymbol{\alpha}_g$ and masses $\pi_g, g = 1, \ldots, G$ of the (discrete) mixing distribution are treated as unknown and estimated together with fixed model parameters, $\boldsymbol{\beta}$ and ϕ. Although the likelihood function in equation (8) can be directly maximized, the EM algorithm (Dempster et al., 1977) represents a common choice for deriving parameter estimates. For this purpose, let us define the binary indicator variable z_{ig} which is equal to one if area i

belongs to the g-th component of the finite mixture. The EM algorithm starts from the definition of the following complete-data log-likelihood:

$$\ell_c(\boldsymbol{\Phi}) = \sum_{i=1}^{m}\sum_{g=1}^{G} z_{ig}\left[\log f_{ig} + \log(\pi_g)\right], \tag{9}$$

where $f_{ig} = f_{y|\alpha}(\boldsymbol{y}_i \mid \boldsymbol{\alpha}_g; \boldsymbol{X}_i)$ is introduced to simplify notation.

The E-step of the algorithm consists in computing the expected value of the complete data log-likelihood (9), conditional on the observed data and the current parameter estimates, say $\hat{\boldsymbol{\Phi}}^{(t)}$. This translates into the derivation of the posterior probability for the i-th small area, $i = 1, \ldots, m$, to belong to the g-th component of the finite mixture, $g = 1, \ldots, G$. Thus, at the generic $(t+1)$-th iteration of the algorithm, we need to derive the following quantities:

$$\tau_{ig}^{(t+1)} = E_{\hat{\boldsymbol{\Phi}}^{(t)}}\left[z_{ig} \mid \boldsymbol{y}_i, \hat{\boldsymbol{\Phi}}^{(t)}\right] = \Pr\left(z_{ig} = 1 \mid \boldsymbol{y}_i, \hat{\boldsymbol{\Phi}}^{(t)}\right) = \frac{\pi_g^{(t)} f_{ig}^{(t)}}{\sum_{l=1}^{G} \pi_l^{(t)} f_{il}^{(t)}}, \tag{10}$$

where $f_{ig}^{(t)}$ denotes the joint conditional distribution of the i-th small area in the g-th component of the finite mixture, with model parameters, $\boldsymbol{\Phi}$, equal to the current updates of the algorithm, $\boldsymbol{\Phi}^{(t)}$. The (conditional) expected value of the complete data log-likelihood is given by

$$Q\left(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}}^{(t)}\right) = \sum_{i=1}^{m}\sum_{g=1}^{G} \tau_{ig}^{(t+1)}\left[\log(f_{ig}) + \log(\pi_g)\right]. \tag{11}$$

In the M-step, parameter estimates are obtained by maximizing (11) with respect to model parameters $\boldsymbol{\Phi}$. A closed form expression is available for the mixture component probabilities:

$$\pi_g^{(t+1)} = \sum_{i=1}^{m} \frac{\tau_{ig}^{(t+1)}}{m}.$$

For the parameters in the regression model, updates depend on the nature of the conditional (component-specific) distribution for Y_{ij}. Generally, the problem reduces to a weighted ML estimation problem for generalized linear models, for which standard Newton-type recursions are usually employed.

The E- and the M-step of the algorithm are iterated until convergence, which can be specified in terms of log-likelihood or parameter values, using appropriate relative or absolute norms. A crucial point in the NPML approach is the choice of the number of mixture components. A simple and often used approach consists in estimating model parameters for a fixed G; this value, in turn, is estimated using a step-by-step procedure built on model selection techniques (see e.g. Böhning, 2003). That is, parameter estimates are computed for varying values of G and the model with the best fit, typically measured by penalized likelihood criteria (such as AIC or BIC), is retained.

3.2 Computing the standard errors for parameter estimates

A disadvantage of the EM algorithm is that it does not directly provide estimates for the covariance matrix of model parameter estimates. However, we may rely on the approach discussed by Louis (1982) and, for practical purposes, on the formula described by Oakes (1999). In this respect, let

$$J(\hat{\boldsymbol{\Phi}}) = -\left[\frac{\partial \ell(\boldsymbol{\Phi})}{\partial \boldsymbol{\Phi} \partial \boldsymbol{\Phi}'}\right]_{\boldsymbol{\Phi}=\hat{\boldsymbol{\Phi}}}$$

denote the observed information matrix. According to Oakes (1999), the following expression holds:

$$J(\hat{\boldsymbol{\Phi}}) = -\left\{\left.\frac{\partial^2 Q(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}})}{\partial \boldsymbol{\Phi} \partial \boldsymbol{\Phi}'}\right|_{\boldsymbol{\Phi}=\hat{\boldsymbol{\Phi}}} + \left.\frac{\partial^2 Q(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}})}{\partial \boldsymbol{\Phi} \partial \hat{\boldsymbol{\Phi}}'}\right|_{\boldsymbol{\Phi}=\hat{\boldsymbol{\Phi}}}\right\}, \tag{12}$$

where the first term denotes the expectation of the complete-data Hessian matrix conditional on the observed data and the parameter estimates. This component can be directly obtained at convergence of the EM algorithm. Explicit formulas for deriving such a term are reported in Appendix A. The second component involved in expression (12) is the first derivative of the conditional expectation of the complete data score (given the observed data) with respect to model parameter estimates (at convergence). For simplicity, this quantity is obtained by computing the numerical derivative of the score functions, where only the weights are considered as functions of parameter estimates.

Once the observed information matrix $J(\hat{\boldsymbol{\Phi}})$ is computed, according to Friedl and Kauermann (2000), the covariance matrix of parameter estimates is obtained via the standard sandwich formula (Huber, 1967; White, 1980):

$$V(\hat{\boldsymbol{\Phi}}) = J(\hat{\boldsymbol{\Phi}})^{-1} V^{\star}(\hat{\boldsymbol{\Phi}}) J(\hat{\boldsymbol{\Phi}})^{-1}, \tag{13}$$

where $V^{\star}(\hat{\boldsymbol{\Phi}}) = \sum_{i=1}^{n} S_i(\hat{\boldsymbol{\Phi}}) S_i(\hat{\boldsymbol{\Phi}})'$ denotes the covariance matrix of the score $S(\boldsymbol{\Phi})$ and $S_i(\hat{\boldsymbol{\Phi}})$ is the individual score vector. Such an approach helps stabilize the estimate of the observed information and ensures robustness to potential model misspecification.

4 Nonparametric EBP and MSE approximation

As in Section 2, let us consider a finite population of size N which can be partitioned into m non-overlapping domains or small areas; we aim at predicting of a (possibly non-linear) function of fixed and random parameters, $\zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\pi})$ where $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_1', \ldots, \boldsymbol{\alpha}_G')'$ and $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_G)'$. Under the proposed approach, the NonParametric Best Predictor (np-BP) of ζ is defined according to

the following expression:

$$\tilde{\zeta}^{\text{np-BP}} = E_{\alpha|y}\left[\zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\pi}) \mid \boldsymbol{y}\right] = \sum_{g=1}^{G} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}_g, \pi_g) \prod_{i=1}^{m} \tau_{ig}, \tag{14}$$

where τ_{ig} is the posterior probability for the i-th small area to belong to the g-th component of the finite mixture. That is, expression (14) denotes the expected value of $\zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\pi})$ with respect to the posterior distribution of the random parameters $\boldsymbol{\alpha}$. Since this latter is a discrete distribution, it is evident that integrals appearing in equation (1) directly translate into summations which are much simpler to be computed, without the use for Monte Carlo approximations or numerical integration techniques.

As for the BP estimator defined in Section 2, an estimate of $\tilde{\zeta}^{\text{np-BP}}$ can be obtained by replacing model parameters $\boldsymbol{\beta}, \boldsymbol{\alpha}$, and $\boldsymbol{\pi}$ by consistent estimates. Here, we consider those derived from the EM algorithm described in Section 3.1. In the following, we will refer to such a quantity as the nonparametric Empirical Best Predictor (np-EBP) of ζ, and denote it by $\hat{\zeta}^{\text{np-EBP}} = \tilde{\zeta}^{\text{np-BP}}(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\pi}})$. To evaluate the quality of predictions obtained using this approach, we follow similar arguments to those detailed in Jiang (2003). In this respect, the MSE of the np-EBP is given by:

$$\text{MSE}(\hat{\zeta}^{\text{np-EBP}}) = d(\boldsymbol{\Phi}) + \frac{1}{m} e(\boldsymbol{\Phi}) + o_p\left(\frac{1}{m}\right), \tag{15}$$

where

$$\begin{aligned} d(\boldsymbol{\Phi}) &= E_\alpha[(\zeta)^2] - E_y[(\tilde{\zeta}^{\text{np-BP}})^2] \\ &= \sum_{g=1}^{G} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}_g, \pi_g)^2 \prod_{i=1}^{m} \Pr(\boldsymbol{\alpha}_i = \boldsymbol{\alpha}_g) - E_y\left[\left(\sum_{g=1}^{G} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}_g, \pi_g) \prod_{i=1}^{m} \tau_{ig}\right)^2\right] \\ &= \sum_{g=1}^{G} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}_g, \pi_g)^2 \, \pi_g^m - E_y\left[\left(\sum_{g=1}^{G} \zeta(\boldsymbol{\beta}, \boldsymbol{\alpha}_g, \pi_g) \prod_{i=1}^{m} \tau_{ig}\right)^2\right]. \end{aligned}$$

The term $e(\boldsymbol{\Phi})$ can be obtained by computing model derivatives with respect to $\boldsymbol{\beta}$, $\boldsymbol{\alpha}_g$, and $\pi_g, g = 1, \ldots, G$, together with the covariance matrix of model parameter estimates defined in equation (13). Here, it is worth noticing that the computational burden required to obtain the above quantities is much lower than that required for the approach in Jiang (2003). Intractable integrals appearing in equations (4)-(5) all translate into simpler summations which can be solved analytically, with a computational complexity which is linear in the integral dimension and remains constant regardless the number of observed small areas m.

With respect to the MSE of the EBP, an estimator of $\text{MSE}(\hat{\zeta}^{\text{np-EBP}})$ is obtained by replacing $\boldsymbol{\Phi}$ in (15) by a consistent estimator, such as that coming from the approach detailed in Section

3. That is,

$$\widehat{\text{MSE}}(\hat{\zeta}^{\text{np-EBP}}) = d(\hat{\boldsymbol{\Phi}}) + \frac{1}{m} e(\hat{\boldsymbol{\Phi}}).$$

As we highlighted before, this approach does not directly lead to an unbiased estimator of $\text{MSE}(\hat{\zeta}^{\text{np-EBP}})$, as, when replacing $\hat{\boldsymbol{\Phi}}$ in $d(\boldsymbol{\Phi})$, we get an error or order $O_p(m^{-1/2})$. However, results described by Boubeta et al. (2016) and Hobza and Morales (2016) showed that the increase in the computational load needed to derive the bias correction term in equation (7) is not rewarded by a real improvement in the quality of the results. For this reason, we decided to overlook this term and simply substitute parameter estimates $\hat{\boldsymbol{\Phi}}$ in the expressions of $d(\boldsymbol{\Phi})$ and $e(\boldsymbol{\Phi})$.

5 A special case: binary data

In this section, we focus on a relevant case in small area estimation. That is, binary responses modeled according to a GLMM with random intercepts. Let Y_{ij} denote the binary response associated to unit j in the i-th small area ($i = 1, \ldots, m, j = 1, \ldots, N_i$), and let α_i denote an area-specific random effect. Last, let $\boldsymbol{x}_{ij}$ denote a p-dimensional vector of covariates, and $\boldsymbol{X}_i$ the matrix of all covariates for the i-th small area. We assume that, conditional on α_i, responses associated to units in small area i are independent Bernoulli random variables with success probability p_{ij}, modeled according to the following mixed logistic model holds:

$$\log \frac{p_{ij}}{1 - p_{ij}} = \eta_{ij} = \alpha_i + \boldsymbol{x}'_{ij}\boldsymbol{\beta}, \tag{16}$$

where $\boldsymbol{\beta}$ is the p-dimensional vector of fixed model parameters that describes the effect of the observed covariates on the logit transform of p_{ij}.

We consider the problem of predicting small area proportions

$$p_i = \frac{1}{N_i} \sum_{j=1}^{N_i} Y_{ij}, \tag{17}$$

and exploit the GLMM in equation (16) to get such predictions. Following the approach detailed in the previous sections, we leave the distribution of the area-specific random effects unspecified and approximate it via a discrete distribution that puts masses $\pi_g > 0$ on locations $\alpha_1, \ldots, \alpha_G$, with $\sum_{g=1}^{m} \pi_g = 1$. These are directly estimated from the data. In this framework, we get the following expression for the success probability for a generic area i belonging to the g-th component of the finite mixture:

$$\log \frac{p_{ijg}}{1 - p_{ijg}} = \eta_{ijg} = \alpha_g + \boldsymbol{x}'_{ij}\boldsymbol{\beta}.$$

Using the standard notation for the Exponential Family and adopting a canonical link function, the joint conditional density for the observed responses coming from the i-th small area is computed as follows:

$$f_{ig} = f_{y|\alpha}(\boldsymbol{y}_i \mid \alpha_g) = \exp\left\{\sum_{j\in s_i}\left[y_{ij}\eta_{ijg} - \log\left(1+e^{\eta_{ijg}}\right)\right]\right\}.$$

Turning back to the problem of estimating small area proportions p_i in equation (17), the corresponding np-BP is given by

$$\begin{aligned}
\tilde{p}_i^{\text{np-BP}} &= \sum_{g=1}^{G} p_{ig}\, \frac{\exp\left[\sum_{j\in s_i} y_{ij}\eta_{ijg} - \sum_{j\in s_i}\log\left(1+e^{\eta_{ijg}}\right)\right]\pi_g}{\sum_{l=1}^{G}\exp\left[\sum_{j\in s_i} y_{ij}\eta_{ijl} - \sum_{j\in s_i}\log\left(1+e^{\eta_{ijl}}\right)\right]\pi_l} \\
&= \sum_{g=1}^{G} p_{ig}\, \frac{\exp\left[\alpha_g y_{i\cdot} - \sum_{j\in s_i}\log\left(1+e^{\eta_{ijg}}\right)\right]\pi_g}{\sum_{l=1}^{G}\exp\left[\alpha_l y_{i\cdot} - \sum_{j\in s_i}\log\left(1+e^{\eta_{ijl}}\right)\right)\pi_l} \\
&= \sum_{g=1}^{G} p_{ig}\, \tau_{igy_{i\cdot}}, \qquad (18)
\end{aligned}$$

where $y_{i\cdot} = \sum_{j\in s_i} y_{ij}$ and $p_{ig} = N_i^{-1}\sum_{j=1}^{N_i} p_{ijg}$.

The NonParametric Empirical Best Predictor (np-EBP) of p_i, that is $\hat{p}_i^{\text{np-EBP}}$, is obtained by substituting ML estimates of model parameters, derived via the EM algorithm detailed in Section 3, into expression (18). That is

$$\hat{p}_i^{\text{np-EBP}} = \sum_{g=1}^{G} \hat{p}_{ig}\,\hat{\tau}_{ig}. \qquad (19)$$

According to equation (15), the quality of predictions obtained via $\hat{p}_i^{\text{np-EBP}}$ can be evaluated through the following MSE expression:

$$\text{MSE}(\hat{p}_i^{\text{np-EBP}}) = E_\alpha[(p_i)^2] - E_y[(\tilde{p}_i^{\text{np-BP}})^2] + E_\alpha[(\hat{p}_i^{\text{np-EBP}} - \tilde{p}_i^{\text{np-BP}})^2], \qquad (20)$$

where

$$E_\alpha[(p_i)^2] = \sum_{g=1}^{G} p_{ig}^2\pi_g$$

and

$$E_y[(\tilde{p}_i^{\text{np-BP}})^2] = \sum_{h=0}^{n_i}(\tilde{p}_{i(h)}^{\text{np-BP}})^2\,\Pr\left(Y_{i\cdot}=h\right).$$

Here, $\tilde{p}_{i(h)}^{\text{np-BP}}$ denotes the np-BP of p_i conditional on $y_{i\cdot} = h$, that is

$$\tilde{p}_{i(h)}^{\text{np-BP}} = \sum_{g=1}^{G} p_{ig} \left[\frac{\exp[\alpha_g h - \sum_{j \in s_i} \log(1 + e^{\eta_{ijg}})]\pi_g}{\sum_{l=1}^{G} \exp[\alpha_l h - \sum_{j \in s_i} \log(1 + e^{\eta_{ijl}})]\pi_l} \right] = \sum_{g=1}^{G} p_{ig} \tau_{ig(h)}.$$

The term $\Pr(Y_{i\cdot} = h)$ is obtained as

$$\Pr(Y_{i\cdot} = h) = \sum_{g=1}^{G} \Pr(Y_{i\cdot} = h \mid \alpha_g)\, \pi_g,$$

where $\Pr(Y_{i\cdot} = h \mid \alpha_g)$ represents the probability of observing h successes in n_i independent, but non identically distributed, Bernoulli trials. This quantity can be obtained from the probability mass function of a Poisson-Binomial random variable (see Chen and Liu, 1997) with parameter $(p_{i1g}, \ldots, p_{in_i g})$.

The last term in equation (20) is obtained as

$$E_\alpha[(\hat{p}_i^{\text{np-EBP}} - \tilde{p}_i^{\text{np-BP}})^2] = \sum_{h=0}^{n_i} \left[\left(\frac{\partial \tilde{p}_{i(h)}^{\text{np-BP}}}{\partial \boldsymbol{\Phi}} \right)' mV(\hat{\boldsymbol{\Phi}}) \left(\frac{\partial \tilde{p}_{i(h)}^{\text{np-BP}}}{\partial \boldsymbol{\Phi}} \right) \right] \Pr(Y_{i\cdot} = h), \tag{21}$$

where $\partial \tilde{p}_{i(h)}^{np-BP} / \partial \boldsymbol{\Phi}$ is the vector of model derivatives conditional on $y_{i\cdot} = h$. Explicit formulas for these derivatives are provided in Appendix B.

A plug-in estimator of $\text{MSE}(\hat{p}_i^{\text{np-EBP}})$, that is $\widehat{\text{MSE}}, (\hat{p}_i^{\text{np-EBP}})$, is obtained by substituting model parameter estimates $\hat{\boldsymbol{\Phi}}$ in the above expressions.

6 Model-based simulations

In this section, we evaluate empirical properties of the proposed approach via a large scale (model-based) simulation study. This consists of $T = 1000$ samples where binary population data are generated under some model assumptions, and sample data are selected from each simulated population. In particular, population data are generated considering $m = 100,\ 200,\ 500$ small areas; then, samples are selected by simple random sampling without replacement within each area. The population and the sample sizes are constant across areas and are fixed at $N_i = 100$ and $n_i = 10$, respectively. According to the simulation study discussed by González-Manteiga et al. (2007), for each unit j in small area i, we generate the target variable $Y_{ij}, i = 1, \ldots, m, j = 1, \ldots, N_i$, from a Bernoulli distribution with success probability defined as

$$p_{ij} = \frac{\exp(\alpha_i + x_{ij}\beta)}{1 + \exp(\alpha_i + x_{ij}\beta)}, \tag{22}$$

with $\beta = 1$ and $x_{ij} \sim \text{Unif}(-1, b_i)$, with $b_i = i/8,\ i/16,\ i/40$ for $m = 100, 200$ and 500, respectively. To evaluate the impact of a parametric assumption on the distribution of the area-specific

random effects, we considered two different scenarios. The first one (Scenario G) uses area-specific random effects generated from a zero mean Gaussian distribution with standard error equal to $\sigma_G = 0.5$. The second scenario (Scenario M) involves area-specific random effects generated from a mixture of Gaussian distributions, obtained as $\alpha_i \sim \delta N(\mu_1, \sigma_M) + (1-\delta)N(\mu_2, \sigma_M)$ where δ represents a random drawn from a Bernoulli distribution $\Pr(\delta = 1) = 0.7$, $\mu_1 = 0$, $\mu_2 = 3$, and $\sigma_M = 0.005$. Based on this latter quantity, it is evident that, under this scenario, the random effect distribution resembles that of a discrete distribution putting masses δ on locations μ_1 and μ_2. The chosen scenarios clearly represents two extreme situations; we expect, in real data applications, the random effect distribution to lie in between them.

In this simulation study, our aim is at comparing the proposed np-EBP of small area proportions in equation (19) with the parametric EBP by Jiang and Lahiri (2001) and the Naive predictor by e.g. González-Manteiga et al. (2007). These latter two are both based on the assumption of Gaussian distributed random effects.

For the EBP, parameter estimates are derived via an approximate ML approach based on a Laplace integration by exploiting the `glmer` function from the R `lme4` package (Bates et al., 2015). Given the estimates, small area proportions and corresponding MSE are derived by adopting the formulas detailed in Section 2. To evaluate the intractable integrals, we follow the approach suggested by Boubeta et al. (2016). That is, we start by generating $B = 2500$ replicates of the area-specific random effects $\alpha_i^{(b)}$ from a Gaussian density with zero mean and variance equal to the corresponding ML estimate. Then we consider their antithetic transform $\alpha_i^{(B+b)} = -\alpha_i^{(b)}$ to obtain $2B$ random effect values. Finally, integrals are approximated by the corresponding empirical means.

In the following, we will denote EBP estimates of small area proportions by $\hat{p}_i^{\text{EBP}}$. Although this approach is optimal, its computational complexity limits its applicability. For this reason, via this simulation study, we aim at understanding whether the np-EBP approach we propose could represent a computationally convenient alternative, that is an optimal choice in terms of minimum MSE and/or a simpler alternative from a computational point of view. For completeness, we also included in the simulation study results from the Naive approach. In this case, parameter estimates are obtained using the PQL approach via the `glmmPQL` function in the R package `MASS` (Venables and Ripley, 2002). Then, to get predictions, parameter estimates are directly plugged in the expression for the area-specific proportions:

$$\hat{p}_i^{\text{Naive}} = \frac{1}{N_i} \sum_{j=1}^{N_i} \frac{\exp(\hat{\alpha}_i + x_{ij}\hat{\beta})}{1 + \exp(\hat{\alpha}_i + x_{ij}\hat{\beta})}. \tag{23}$$

The performance of the small area estimators were evaluated by computing, for each area

Figure 1: Scenario G: Distribution of the absolute bias over areas for $\hat{p}_i^{\text{np-EBP}}$, $\hat{p}_i^{\text{EBP}}$, and $\hat{p}_i^{\text{Naive}}$, for $m = 100$ (left panel), $m = 200$ (central panel), and $m = 500$ (right panel).

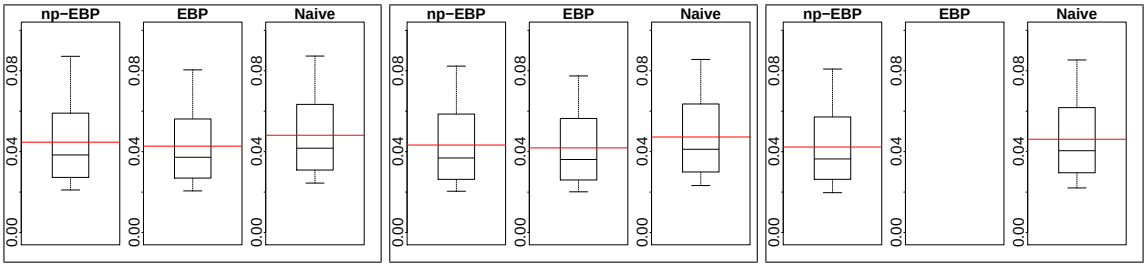

Figure 2: Scenario G: Distribution of the RMSE over areas of $\hat{p}_i^{\text{np-EBP}}$, $\hat{p}_i^{\text{EBP}}$, and $\hat{p}_i^{\text{Naive}}$, for $m = 100$ (left panel), $m = 200$ (central panel), and $m = 500$ (right panel).

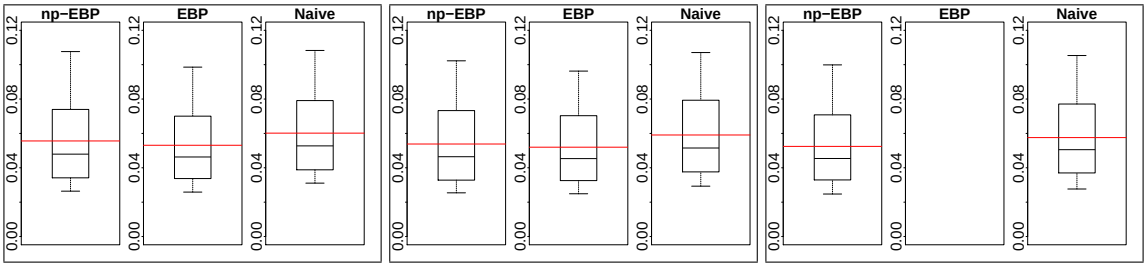

$i = 1, \ldots, m$ the absolute bias (AB) and the root mean squared error (RMSE):

$$\text{AB}_i = T^{-1} \sum_{t=1}^{T} |\hat{p}_{it}^{\text{Model}} - p_{it}|, \quad i = 1, \ldots, m,$$

$$\text{RMSE}_i = \sqrt{T^{-1} \sum_{t=1}^{T} (\hat{p}_{it}^{\text{Model}} - p_{it})^2}, \quad i = 1, \ldots, m,$$

where $\hat{p}_{it}^{\text{Model}}$ denotes the model-based estimated proportion for the i-th small area in the t-th simulated sample, obtained via either the EBP ($\hat{p}_i^{\text{EBP}}$), the np-EBP ($\hat{p}_i^{\text{np-EBP}}$) or the Naive ($\hat{p}_i^{\text{Naive}}$) approach. Figures 1 and 2 show the absolute bias and the actual RMSE distribution across small areas for the np-EBP, the EBP, and the Naive estimators, for $m = 100, 200, 500$, respectively, under Scenario G. As expected, when looking at the first two panels (i.e. $m = 100,\ 200$), the np-EBP performs better than the Naive estimator and slightly worse than the EBP, with a gap that reduces as m increases, both in terms of AB and RMSE. When $m = 500$, performance values of EBP are not showed due to the computational burden required to get the estimates: for one replication, we needed 161.612 minutes on an Intel(R) I5-3330 architecture – 3.0 GHz, and, therefore, we couldn't obtain results in a reasonable amount of time.

Figures 3 and 4 show the performance of the estimators performance under Scenario M. As before, results for the EBP approach when $m = 500$ are not showed for similar computational

Figure 3: Scenario M: Distribution of the absolute bias over areas of $\hat{p}_i^{\text{Naive}}$, $\hat{p}_i^{\text{EBP}}$, and $\hat{p}_i^{\text{np-EBP}}$ for $m = 100$ (left panel), $m = 200$ (central panel), and $m = 500$ (right panel).

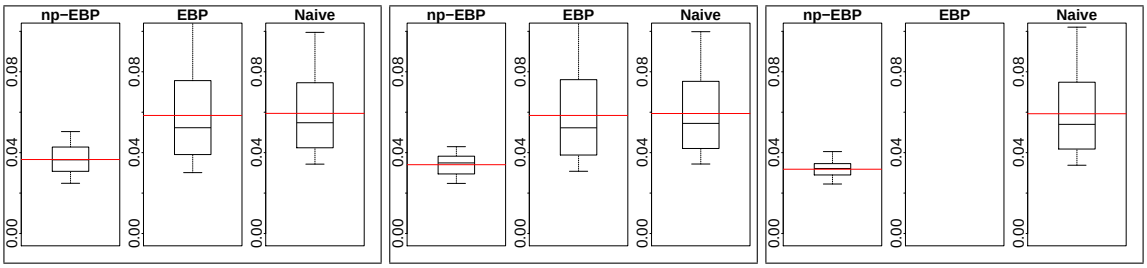

Figure 4: Scenario M: Distribution of the RMSE over areas of $\hat{p}_i^{\text{Naive}}$, $\hat{p}_i^{\text{EBP}}$, and $\hat{p}_i^{\text{np-EBP}}$ for $m = 100$ (left panel), $m = 200$ (central panel), and $m = 500$ (right panel).

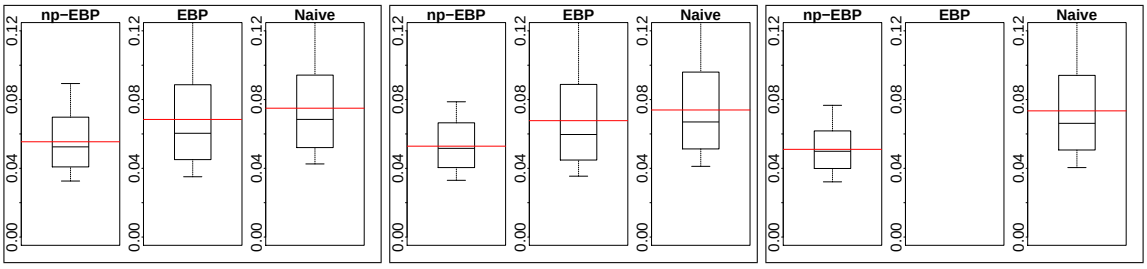

issues. As it is clear by looking at these plots, when the assumption of Gaussian random effects does not hold, parametric approaches seem to produce predictions with a reduced quality when compared to the nonparametric alternative. In particular, we may notice that $\hat{p}_i^{\text{np-EBP}}$ outperforms the two competitors in terms of AB and RMSE. The results of both $\hat{p}_i^{\text{Naive}}$ and $\hat{p}_i^{\text{EBP}}$ worsen slightly as m increases. This may be possibly due to the higher information available and the stronger impact of the random parameter distribution on the overall response variability when the number of observed small areas increases.

We now investigate the performance of MSE estimators to evaluate the accuracy of the predictions discussed so far. In particular, for $\hat{p}_i^{\text{np-EBP}}$, we considered the estimator provided in equation (20); in the following, we will refer to the square root of such a quantity as $\widehat{\text{RMSE}}(\hat{p}_i^{\text{np-EBP}})$. For the estimator $\hat{p}_i^{\text{EBP}}$, we used the approximate MSE estimator proposed by Jiang and Lahiri (2001); in the following, the corresponding square root will be denoted by $\widehat{\text{RMSE}}(\hat{p}_i^{\text{EBP}})$. Last, for the Naive predictor $\hat{p}_i^{\text{Naive}}$, we considered the approach suggested by González-Manteiga et al. (2007), based on linearizing the GLMM in equation (22) and, then, applying the Prasad-Rao MSE approximation for the corresponding linear mixed model; in the following, the square root of such an estimator will be denoted by $\widehat{\text{RMSE}}(\hat{p}_i^{\text{Naive}})$.

Performances of the RMSE estimators were evaluated by considering the ratio between the estimated RMSE of the model-based estimates and the corresponding actual Monte Carlo RMSE

of each small area prediction (R), that is:

$$\mathrm{R}_i = \frac{\sum_{t=1}^{T} \widehat{\mathrm{RMSE}}(\hat{p}_{it}^{\mathrm{Model}})}{\sqrt{\sum_{t=1}^{T} (\hat{p}_{it}^{\mathrm{Model}} - p_{it})^2}}.$$

The distribution over areas of the ratios for Scenario G and Scenario M is reported in Figure 5. The RMSE estimator of np-EBP seems to perform generally better than the alternatives, with precision that increases with m. Under Scenario G, the RMSE estimator of the EBP seems to underestimate the actual Monte Carlo RMSE, with a ratio which is always lower than 1 when $m = 100$. Again, the quality of results improves with m, even though it is always lower than that provided by the proposed approach. Such a finding my be possibly due to the estimation of the covariance matrix of parameter estimates which is not as accurate as expected when $B^* = 250$ bootstrap samples are considered. In fact, with this regard it is worth noticing that Boubeta et al. (2016) highlighted the need of a very accurate estimate of the covariance matrix of parameter estimates to ensure the quality of results. For this reason, in their simulation study, the authors suggested to estimate $V(\hat{\mathbf{\Phi}})$ by running a Monte Carlo experiment based on 10^4 iterations in advance. In practice, when dealing with large sample sizes, such an approach is computationally very expensive and this is the reason why we have considered a bootstrap approach based on 250 iterations only. On the other hand, the estimator suggested by González-Manteiga et al. (2007) seems to overestimate the actual Monte Carlo RMSE in all the scenarios we considered in this simulation study.

When dealing with non-Gaussian random effects, the MSE estimator of the np-EBP has again a good performance with average ratio over areas close to 1 for all values of m. Under this scenario, also the MSE estimator of the EBP works quite well, whereas that for the Naive consistently underestimates the actual RMSE.

Finally, Table 1 shows the mean coverage rate (CR) of nominal 95% Gaussian-type confidence intervals over simulations, that is

$$\mathrm{CR}_i = T^{-1} \sum_{t=1}^{T} \mathbb{1}\left(|\hat{p}_{it} - p_{it}| \leqslant 1.96 \times \widehat{\mathrm{RMSE}}(p_{it}^{\mathrm{Model}}) \right), \quad i = 1, \ldots, m,$$

and the computational time required to get estimates of the actual RMSE. The proposed np-EBP has a good performance with average empirical coverage of approximately 93% in all cases but with Scenario G and when $m = 100$. On the other hand, both EBP and, particularly, Naive show a more erratic behavior. The former shows a relatively lower coverage rate under Scenario G and over coverage under Scenario M. This behavior is reversed for Naive.

To conclude, we also offer a comparison between MSE estimators in terms of computational complexity. The last columns of Tables 1 reports the computational time (averaged over simulations) required to get the estimates on an Intel(R) I5-3330 architecture (3.0 GHz) under each

Figure 5: Distribution of the RMSE ratio over areas for the np-EBP, the EBP, and the Naive approach, for $m = 100$ (left panel), $m = 200$ (central panel), and $m = 500$ (right panel), under Scenario G (upper panel) and Scenario M (lower panel).

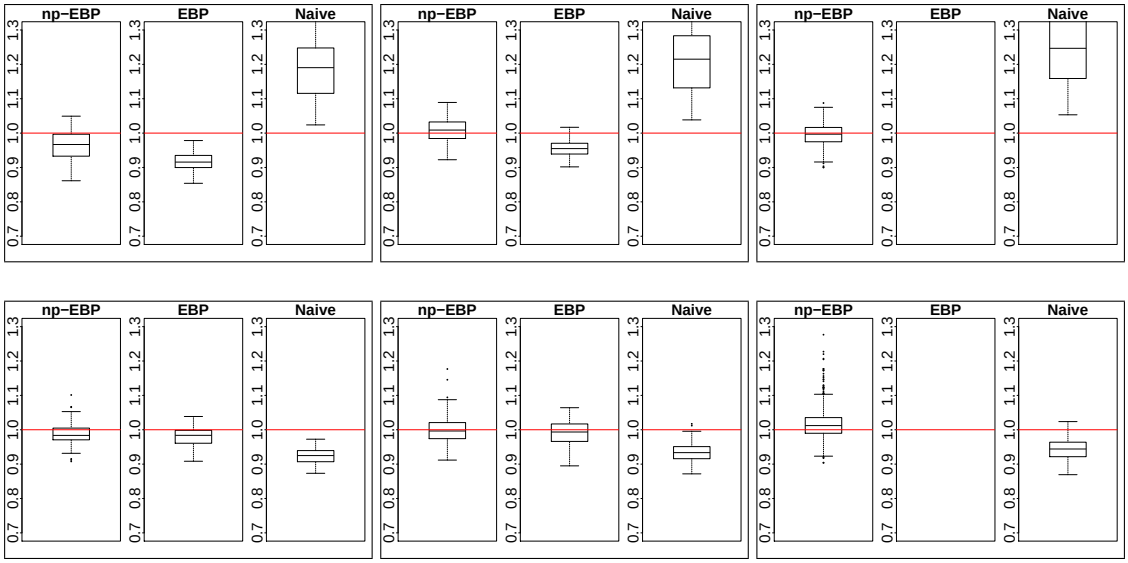

simulation setting. As it can be seen, $\widehat{\text{RMSE}}(\hat{p}_i^{\text{np-EBP}})$ shows good computational performance when compared to the Naive and the EBP estimators. Obviously, the computational load increases with m, but it looks in all cases reasonable. It is interesting to notice that the average time for the Naive MSE estimator is low for $m = 100, 200$, but it increases exponentially with m, even though it is always lower than np-EBP. Last, the MSE estimator for the EBP by Jiang and Lahiri (2001) needs an average time of estimation that becomes huge when $m = 500$ and this makes this approach unfeasible for empirical applications with large m, as for the one we will discuss in Section 7.

7 Estimating unemployment incidence for Local Labor Market Areas in Italy

In this section, we use data from the first quarter of the 2012 edition of the Italian Labor Force Survey (ILFS) and employ the proposed np-EBP to estimate unemployment incidence for 611 Local Labor Market Areas (LLMAs) in Italy. LLMAs are unplanned domains obtained as clusters of municipalities cutting across provinces (LAU1) which are ILFS finest planned domains. In 2011, with the last Census, LLMAs have been re-defined in terms of daily working commuting flows. The ILFS, for each quarter, collects information on almost $70,000$ households in $1,246$ municipalities for a total of $175,000$ individuals. The ILFS provides quarterly estimates of the main aggregates of labor market such as employment status, type of work and work

Table 1: Average coverage rate over areas and computational time (in minutes) of $\widehat{\text{RMSE}}(\hat{p}_i^{\text{Naive}})$, $\widehat{\text{RMSE}}(\hat{p}_i^{\text{EBP}})$, and $\widehat{\text{RMSE}}(\hat{p}_i^{\text{np-EBP}})$, for $m = 100, 200, 500$.

	Coverage			Computational Time		
	$m = 100$	$m = 200$	$m = 500$	$m = 100$	$m = 200$	$m = 500$
`Scenario` G						
np-EBP	0.879	0.920	0.939	0.078	0.124	0.310
EBP	0.864	0.912	-	16.907	41.385	-
Naive	0.962	0.967	0.973	0.005	0.018	0.206
`Scenario` M						
np-EBP	0.927	0.930	0.930	0.079	0.159	0.429
EBP	0.966	0.974	-	17.229	42.112	-
Naive	0.903	0.906	0.910	0.004	0.018	0.219

experience, stratified by sex, age and region.

Direct survey estimates can not be computed and/or published for most of the LLMAs because 158 out of 611 are out of sample areas and many LLMAs have a small sample size that provides estimates with unacceptable large coefficient of variation. However, for local administrators, it is important to get information on the local labor market which can be used to develop, manage, evaluate, and report on labor market policies. The sample size and the number of small areas are particularly large in this application. According to the simulation results in Section 6, the np-EBP is a possible approach to be used as (i) it performs better than the Naive predictor in terms of bias and efficiency; (ii) the computational complexity of the MSE estimator for the parametric EBP makes it inadequate to be used when dealing with a large number of small areas and/or large sample sizes, as those from the ILFS dataset.

7.1 Preliminary analysis

Sample data is available on $93,217$ units in 453 out of 611 LLMAs from the the first quarter of 2012. This time span is the closest to the date of the 2011 Census, so that auxiliary information on the population coming from this source can be usefully employed in the small area model. In particular, we are interested in the *employment status* variable. This can take one out of three different categories: employed (53.6%), unemployed (6.6%), and inactive (39.8%), where these are sample percentages. Together with this response, a number of explanatory variables is available at the population level. In particular, we focus on the following:

- *Sex-Age*: a categorical variable with six categories corresponding to female or male (F/M) and three age groups (15-24, 25-34, and 35-65);

Table 2: Sample percentage distribution of *Unemployed status* by *Sex-Age*

	Unemployed	Employed	Inactive
M: 15-24	11.0	21.4	67.6
M: 25-34	11.7	71.6	16.7
M: 35-65	5.3	70.8	23.9
F: 15-24	8.8	14.1	77.1
F: 25-34	11.6	53.9	34.5
F: 35-65	4.2	48.3	47.5

Table 3: Sample percentage distribution of *Unemployed status* by *Educationial level*

	Unemployed	Employed	Inactive
Primary school or less	4.6	24.9	70.5
Middle school	6.9	44.5	48.6
High school	7.2	62.8	30.0
University degree or beyond	5.6	75.7	18.7

- *Educational level*: a categorical variable with four categories corresponding to: no education or primary school diploma, secondary school diploma, high school diploma, university degree or beyond;

- *U-count*: number of unemployed in a given sex-age group for each LLMA according to the 2011 Census.

Such covariates were also used in other studies on small area labor force characteristics (see e.g. Molina et al., 2007; Chambers et al., 2016).

To have a first insight of the data, we report in Tables 2-3 the sample distribution of the *employment status* stratified by *Sex-Age* and *Educational level*, respectively. Unemployment incidence is generally higher for people aged 25-34, regardless of the sex. For the other age classes, unemployment is less frequent among females. However, by looking at the last column of the table, we may notice that females are more frequently *inactive*, regardless the age group, when compared to males. This is likely due to their engagement in housekeeping and explains why unemployment incidence is lower within this group. Similarly, by focusing on Table 3, we may observe that the percentage of unemployment is relatively larger for more educated individuals. However, also in this case, by looking at the last column, it is evident that such a finding is mainly related to the job search. For instance, 70% of individuals with a primary school diploma or less are out of the job market, and this can be explained by their relatively older age.

In the next sections, we use these auxiliary variables to build a mixed logistic model with discrete, area-specific, random effects and provide predictions for small area unemployment

incidences. In particular, the response Y_{ij} takes value 1 if unit j in small area i is unemployed and 0 otherwise.

7.2 The model

We considered different model specifications and a varying number of mixture components ($G = 2, \ldots, 6$) for the random effects distribution. The optimal solution, corresponding to the smallest AIC value, is based on $G = 3$ components and includes in the linear predictor *Sex-Age*, *Educational level* and log(*U-count*). Model parameter estimates, standard errors and p-values are set out in Table 4. Looking at the estimates for *Sex-Age*, we may notice that, when controlling for the effect of other explanatory variables in the model and for the effect of unobserved heterogeneity, the odd of being unemployed for younger people is higher than that of individuals aged $35 - 65$ years. For instance, the odds of being unemployed for males who belong to the category $25 - 34$ are 1.125 times those of males aged $16 - 24$; on the other hand, for males who are aged $35 - 65$ years, the odds are 0.459, that is, 54.1% lower than those for the baseline category. Such differences are even stronger for females.

As regards the *Educational level*, the odds of being unemployed for high educated people are higher than those for other categories (parameter estimates for middle and high school diplomas are all positive). Finally, having a University degree or higher education does not seem to have a significant effect with respect to the baseline. These findings are in line with the results reported in the preliminary analysis: low educated female and relatively younger individuals (the reference category) are more frequent in the inactive category. Last, as expected, results reported in Table 4 suggest that the probability of being unemployed increases as the total number of registered unemployed at the 2011 census increases.

Figure 6 shows the estimated prior and posterior distribution for random effects. The latter is obtained by computing the posterior mean of the area-specific random intercepts as

$$\alpha_i = \sum_{g=1}^{G} (\alpha_g - \bar{\alpha}) \tau_{ig},$$

where $\bar{\alpha} = \sum_g \alpha_g \pi_g$ is the overall intercept reported in Table 4. This distribution is clearly skewed and the Gaussian assumption, which is typical of standard generalized linear mixed models, seems not to be adequate in this case. This finding provides evidence that our approach is a useful tool in the analysis of these data.

7.3 Small area predictions

As highlighted in Section 2, in general, to predict the target variable, we need the covariate values, $\boldsymbol{x}_{ij}$, to be known for all units in the population. This would require access e.g. to census

Table 4: Parameter estimates, standard errors and corresponding p-values for the mixed logistic model fitted to the ILFS data. $G = 3$ number of components of the mixture for the random effects distribution.

	Estimate	SE	p-value
Intercept	-3.052	0.176	0.038
M: 25-34	0.118	0.057	< 0.001
M: 35-65	-0.787	0.054	< 0.001
F: 15-24	-0.222	0.062	0.023
F: 25-34	0.113	0.050	< 0.001
F: 35-65	-1.039	0.052	< 0.001
Middle school	0.206	0.059	< 0.001
High school	0.235	0.062	< 0.001
University degree or beyond	-0.032	0.083	0.682
$\log(U\text{-}count)$	0.113	0.012	< 0.001

micro-data. However, in the important and special case where the components of $\boldsymbol{x}_{ij}$ are all categorical, or take on a finite number of values, as in this application, we may notice that using the method described in this paper only requires the corresponding area level cross-tabulations to be available.

Figure 7 shows the map of unemployment incidence prediction for the 611 LLMAs obtained using direct estimation, parametric EBP, and the proposed np-EBP. The patterns of unemployment produced by the proposed approach are consistent with those obtained by both the direct approach and the EBP method. As expected, model-based maps are smoother when compared to direct estimates and relatively larger values for unemployment incidences are mainly located in the South of Italy and in the Islands.

To assess the quality of predictions, we used a set of diagnostic tools based on the requirement that model-based small area estimates should be coherent with, in the sense of being close to, the corresponding unbiased direct estimates, albeit more precise. Figure 8 shows the estimates derived from the np-EBP approach versus the Direct, the EBP and the Naive estimates, respectively. From this figure (first panel), we may observe that our approach leads to predictions which are very similar to those provided by a direct approach, with a correlation coefficient equal to 0.881. From the remaining panels in Figure 8, it is evident that model-based estimates for unemployment incidence are all very close to each other, with correlation coefficients equal to 0.978 (np-EBP vs. EBP) and to 0.977 (np-EBP vs. Naive).

Coherence of Direct and np-EBP estimates can be also evaluated by computing a goodness-

Figure 6: Estimated prior (a) and posterior (b) distribution for the random effects.

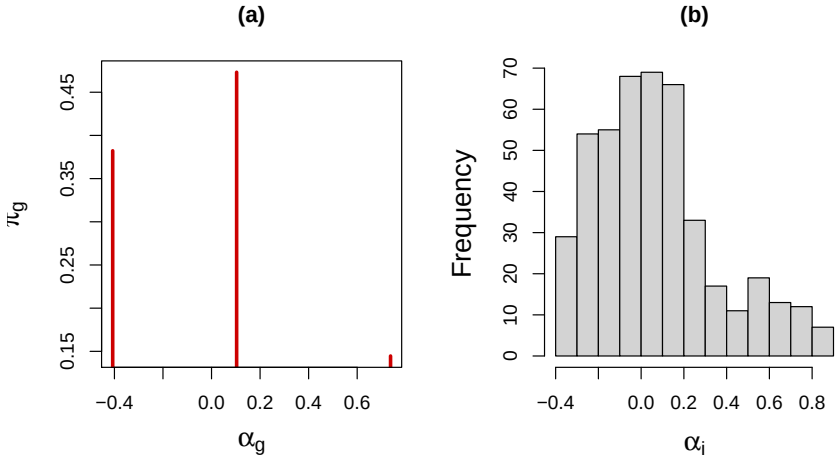

of-fit diagnostic (Brown et al., 2001), which is obtained from the following Wald-test statistic:

$$W = \sum_{i=1}^{m} \frac{(\hat{p}_i^{\text{Direct}} - \hat{p}_i^{\text{np-EBP}})^2}{\widehat{Var}(\hat{p}_i^{\text{Direct}}) + \widehat{MSE}(\hat{p}_i^{\text{np-EBP}})}, \tag{24}$$

where the estimated MSE of the np-EBP proportions is calculated by using formulas in Section 5. The above test is based on the idea that, should model-based estimates be close to the "true" small area parameters of interest, the unbiased direct estimates could be considered as random variables with expected value equal to the value of the corresponding model-based estimate. Here, $W = 360.56$ and such a value needs to be compared to the 95-th percentile of a χ^2 distribution with 452 d.f., i.e. $\chi^2_{452,0.95} = 502.56$. In this respect, we may conclude that model-based estimates are not significantly different from the direct estimates.

To assess the potential gain in precision we have when using model-based estimates based on the np-EBP approach instead of Direct estimates, we compared in Figure 9 the empirical cumulative density functions (ecdfs) of the Coefficients of Variation associated to such predictions. The first panel uses CV's from all areas, while the second (third) one focuses on small areas with a sample size smaller (higher) than 100. As it is clear, by looking at the first panel, the ecdf corresponding to np-EBP almost always dominates the one for the Direct estimates, highlighting how CV values for the former approach are lower than those estimated with the latter. Only for very small CV values, the ecdfs show an inverse relation: CV's for Direct predictions are smaller than those of the np-EBP ones. This is more evident from the second panel. This difference is related to the presence of some areas with a small sample size for which $\widehat{RMSE}(\hat{p}_i^{\text{Direct}})$ is zero or is very close to zero. When we move to higher CVs, np-EBP always provides much smaller

Figure 7: Maps of the estimated unemployment incidences for LLMAs in Italy in 2012: direct estimates, EBP and np-EBP estimates.

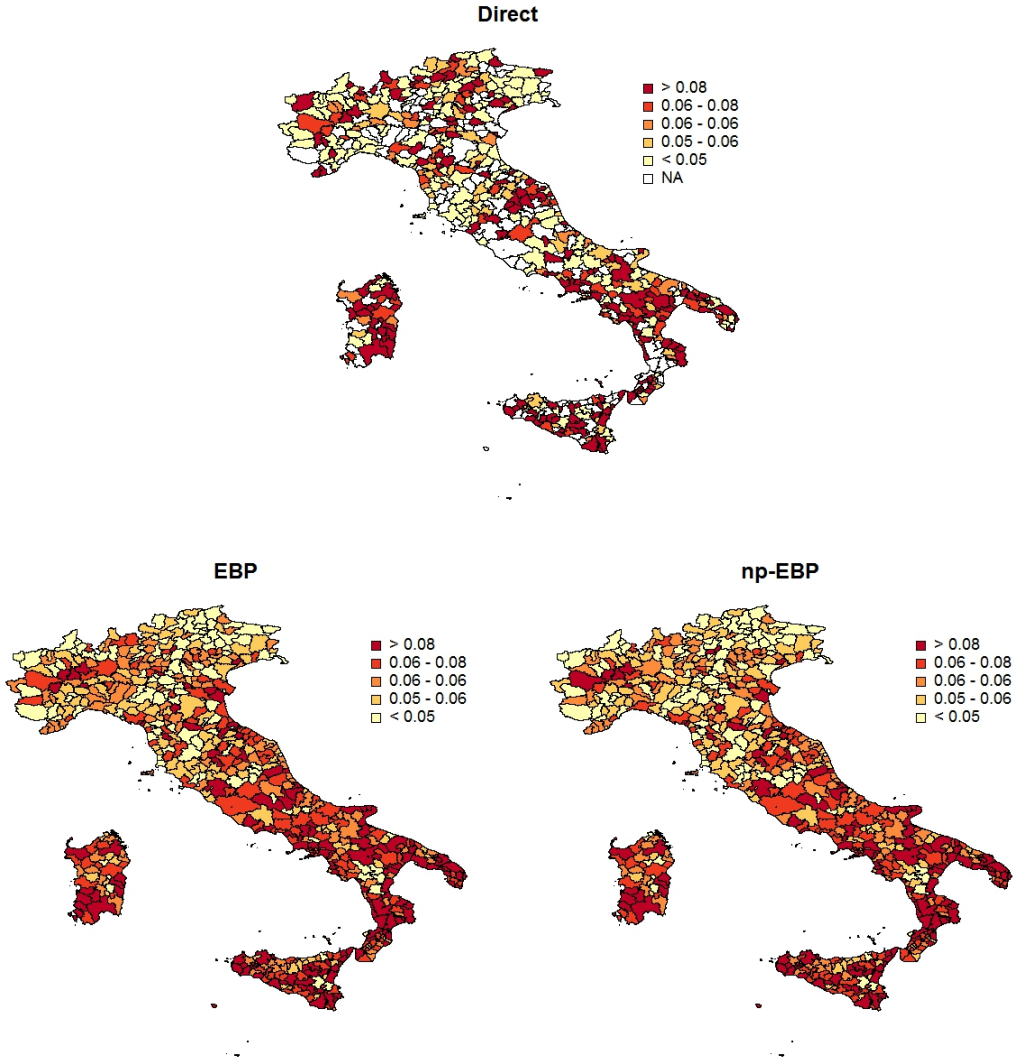

Figure 8: np-EBP estimates of small area proportions versus the corresponding Direct (left), EBP (centre), and Naive estimates (right).

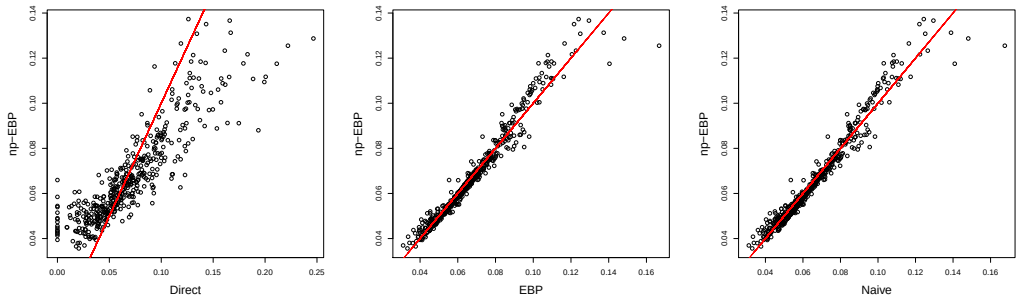

Figure 9: CV's empirical cumulative density functions for the np-EBP and the direct estimator.

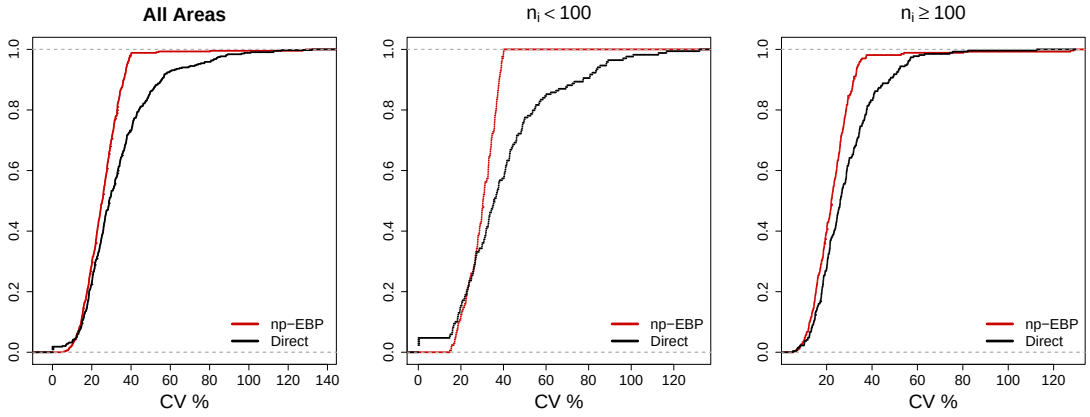

CV values than the Direct approach and they are always smaller than 40%. When focusing on the third panel in Figure 9, it is evident that, when the sample size gets larger, direct estimates and model-based estimates have more similar CV values, although model-based estimates are still consistently smaller.

8 Conclusions

The paper offers some tools to derive best prediction in the presence of data with clustered/hierarchical structure for general responses with distribution in the Exponential family. In particular, we proposed a nonparametric version of EBP and the corresponding MSE approximation using a NonParametric Maximum Likelihood approach (Aitkin, 1996, 1999). Motivated from the real data application to a dataset on unemployment incidence in LLMAs in Italy, we then focused on a dichotomous response modelled according to a generalized linear mixed model with random intercepts. This is a relevant case in the small area estimation framework.

Simulation experiments showed that the proposed predictor performs equally or better than the competitors. In particular, when the assumption of Gaussian distributed random effects does not hold, the np-EBP performs better than the corresponding parametric version, based on a Gaussian assumption. Empirical experiments highlighted the need of a much lower computational time to get both predictions and, above all, the corresponding MSE, when compared to the parametric alternative by Jiang (2003). In particular, the simulation study, in which different sample sizes were considered, showed that the proposed approach is quite appropriate, especially for large m. Such a gain comes from the discrete nature of the mixing distribution which substantially simplifies calculations needed to get EBPs and the corresponding MSEs. The use of the proposed methodology is facilitated by the availability of a (computationally efficient algorithm) developed in `R` and available from the authors upon request.

We also illustrated the benefits of the proposed approach by applying it to the estimation of unemployment incidence for LLMAs in Italy in 2012. In this context, direct estimates can not be published for most of the LLMAs given the unacceptable large value of the coefficient of variation for those areas with a small sample size. In this respect, model-based approaches represent an unavoidable strategy to be used. However, the sample size and the number of small areas are particularly large in this application and this makes implementation of the EBP particularly cumbersome, and the evaluation of its precision prohibitive. Results from this application indicated that the proposed methodology leads to estimates which are coherent, but more efficient than, direct estimates, and that are comparable with alternative model-based estimates.

Although the proposed approach is presented for general responses in the Exponential Family, we have not explored the behavior of the corresponding small area np-EBPs for counts or multinomial responses. A possible extension to multi-category outcomes is however quite straightforward in the current class of models. Finally, suitable extensions of the proposed approach to allow for spatial correlation of the small areas could be envisioned by properly modeling prior mixture probabilities as a function of the mixture component of the neighboring areas.

Acknowledgements

The work of Marino, Ranalli and Salvati has been developed under the support of the project PRIN-SURWEY (grant 2012F42NS8, Italy). The work of Alfó has been supported by the project "Mixture and latent variable models for causal inference and analysis of socio-economic data" (grant RBFR12SHVV, FIRB - Futuro in Ricerca).

References

Aitkin, M. (1996). A general maximum likelihood analysis of overdispersion in generalized linear models. *Statistics and Computing*, 6:251–262.

Aitkin, M. (1999). A general maximum likelihood analysis of variance components in generalized linear models. *Biometrics*, 55:117–128.

Aitkin, M. and Alfò, M. (2003). Longitudinal analysis of repeated binary data using autoregressive and random effect modelling. *Statistical Modelling*, 3:291–303.

Alfò, M. and Aitkin, M. (2006). Variance component models for longitudinal count data with baseline information: epilepsy data revisited. *Statistics and Computing*, 16:231–238.

Alfò, M., Salvati, N., and Ranalli, M. G. (2017). Finite mixtures of quantile and M-quantile regression models. *Statistics and Computing*, 27:547–570.

Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67:1–48.

Böhning, D. (1982). Convergence of Simar's algorithm for finding the maximum likelihood estimate of a compound poisson process. *The Annals of Statistics*, 10:1006–1008.

Böhning, D. (2003). The EM algorithm with gradient function update for discrete mixtures with known (fixed) number of components. *Statistics and Computing*, 13:257–265.

Boubeta, M., Lombardía, M. J., and Morales, D. (2016). Empirical best prediction under area-level Poisson mixed models. *Test*, 25(3):548–569.

Breslow, N. E. and Clayton, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association*, 88:9–25.

Brown, G., Chambers, R., Heady, P., and Heasman, D. (2001). *Evaluation of small area estimation methods: an application to unemployment estimates from the UK LFS.* In Proc. Statistics Canada Symp. Achieving Data Quality in a Statistical Agency: a Methodological Perspective. Hull: Statistics Canada.

Chambers, R., Salvati, N., and Tzavidis, N. (2016). Semiparametric small area estimation for binary outcomes with application to unemployment estimation for local authorities in the UK. *Journal of the Royal Statistical Society: Series A*, 179:453–479.

Chen, S. X. and Liu, J. S. (1997). Statistical applications of the Poisson-binomial and conditional Bernoulli distributions. *Statistica Sinica*, 7:875–892.

Dempster, A., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B*, 39:1–38.

Friedl, H. and Kauermann, G. (2000). Standard errors for EM estimates in generalized linear models with random effects. *Biometrics*, 56:761–767.

González-Manteiga, W., Lombardía, M., Molina, I., Morales, D., and Santamaría, L. (2007). Estimation of the mean squared error of predictors of small area linear parameters under a logistic mixed model. *Computational Statistics and Data Analysis*, 51:2720–2733.

Hinde, J. and Wood, A. (1987). Binomial variance component models with a non-parametric assumption concerning random effects. In *Longitudinal Data Analysis*, pages 110–127. Avebury, Aldershot, Hants.

Hobza, T. and Morales, D. (2016). Empirical best prediction under unit-level logit mixed models. *Journal of Official Statistics*, 32:661–692.

Huber, P. (1967). The behaviour of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the 5th Berkeley Symposium*, volume 1, pages 221–233.

Jiang, J. (1998). Consistent estimators in generalized linear mixed models. *Journal of the American Statistical Association*, 93:720–729.

Jiang, J. (2003). Empirical best prediction for small-area inference based on generalized linear mixed models. *Journal of Statistical Planning and Inference*, 111:117–127.

Jiang, J. and Lahiri, P. (2001). Empirical best prediction for small area inference with binary data. *Annals of the Institute of Statistical Mathematics*, 53:217–243.

Kiefer, J. and Wolfowitz, J. (1956). Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics*, 27:887–906.

Krishnakumar, J. (2006). Chapter 5 time invariant variables and panel data models: A generalised Frisch-Waugh theorem and its implications. In Baltagi, B. H., editor, *Panel Data EconometricsTheoretical Contributions and Empirical Applications*, volume 274 of *Contributions to Economic Analysis*, pages 119 – 132. Elsevier.

Laird, N. (1978). Nonparametric maximum likelihood estimation of a mixing distribution. *Journal of the American Statistical Association*, 73:805–811.

Laird, N. M. and Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, 38:963–974.

Lindsay, B. G. (1983a). The geometry of mixture likelihoods: a general theory. *Ann. Statist.*, 11:86–94.

Lindsay, B. G. (1983b). The geometry of mixture likelihoods, Part II: the exponential family. *The Annals of Statistics*, 11:783–792.

Louis, T. A. (1982). Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44:226–233.

McCulloch, C. E. (1997). Maximum likelihood algorithms for generalized linear mixed models. *Journal of the American statistical Association*, 92:162–170.

Molina, I., Saei, A., and Lombardía, M. (2007). Small area estimates of labour force participation under a multinomial logit mixed model. *Journal of the Royal Statistical Society: Series A*, 70:265–283.

Neuhaus, J. M. and McCulloch, C. E. (2006). Separating between-and within-cluster covariate effects by using conditional and partitioning methods. *Journal of the Royal Statistical Society: Series B*, 68:859–872.

Oakes, D. (1999). Direct calculation of the information matrix via the em. *Journal of the Royal Statistical Society: Series B*, 61:479–482.

Opsomer, J. D., Claeskens, G., Ranalli, M. G., Kauermann, G., and Breidt, F. J. (2008). Nonparametric small area estimation using penalized spline regression. *Journal of the Royal Statistical Society: Series B*, 70:265–283.

Pinheiro, J. and Bates, D. (1995). Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics*, 4:12–35.

Prasad, N. and Rao, J. (1990). The estimation of mean squared error of small-area estimators. *Journal of the American Statistical Association*, 85:163–171.

Rao, J. N. and Molina, I. (2015). *Small area estimation*. John Wiley & Sons.

Saei, A. and Chambers, R. (2003). Small area estimation under linear and generalized linear mixed models with time and area effects. In *S3RI Methodology Working Papers*, pages 1–35. Southampton Statistical Sciences Research Institute, Southampton.

Simar, L. (1976). Maximum likelihood estimation of a compound poisson process. *The Annals of Statistics*, 4:1200–1209.

Venables, W. N. and Ripley, B. D. (2002). *Modern Applied Statistics with S*. Springer, New York, fourth edition.

Wedel, M., DeSarbo, W. S., Bult, J. R., and Ramaswamy, V. (1993). A latent class Poisson regression model for heterogeneous count data. *Journal of Applied Econometrics*, 8:397–411.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and direct test for heteroskedasticity. *Econometrica*, 48:817–838.

A Posterior expectation of the complete-data Hessian

In this section, we provide explicit formulas for deriving the expected value of the complete data Hessian matrix, conditional on the observed data and the parameter estimates. Let us denote by $h[\cdot] = g^{-1}[\cdot]$ the inverse link function. As a first step, we need to compute the score functions associated with each element of $\boldsymbol{\Phi}$, based on the sample data only; these are given by

$$
\begin{aligned}
S(\boldsymbol{\alpha}_g) &= \frac{\partial Q(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}})}{\partial \boldsymbol{\alpha}_g} = \sum_{i=1}^{m} \tau_{ig} S_{ig}(\boldsymbol{\alpha}_g) = \sum_{i=1}^{m} \tau_{ig} \sum_{j \in s_i} \left\{ y_{ij} - h[\eta_{ijg}] \right\} \boldsymbol{w}_{ij}, \quad g = 1, \ldots, G, \\
S(\boldsymbol{\beta}) &= \frac{\partial Q(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}})}{\partial \boldsymbol{\beta}} = \sum_{i=1}^{m} \sum_{g=1}^{G} \tau_{ig} S_{ig}(\boldsymbol{\beta}) = \sum_{i=1}^{m} \sum_{g=1}^{G} \tau_{ig} \sum_{j \in s_i} \left\{ y_{ij} - h[\eta_{ijg}] \right\} \boldsymbol{x}_{ij}, \\
S(\pi_g) &= \frac{\partial Q(\boldsymbol{\Phi} \mid \hat{\boldsymbol{\Phi}})}{\partial \pi_g} = \sum_{i=1}^{m} \left[\frac{\tau_{ig}}{\pi_g} - \frac{\tau_{iG}}{\pi_G} \right], \quad g = 1, \ldots, G-1,
\end{aligned} \tag{25}
$$

where τ_{ig} represents the posterior probability for the i-th small area to belong to the g-th component of the finite mixture. Given the score functions above, the expectation of the complete

data Hessian matrix can be obtained from the following equations:

$$H(\boldsymbol{\alpha}_g, \boldsymbol{\alpha}_g) = \frac{\partial S(\boldsymbol{\alpha}_g)}{\partial \boldsymbol{\alpha}_g'} = -\sum_{i=1}^{m} \tau_{ig} \sum_{j \in s_i} \left\{ \frac{\partial h[\eta_{ijg}]}{\partial \eta_{ijg}} \right\} \boldsymbol{w}_{ij} \boldsymbol{w}_{ij}', \quad g = 1, \ldots, G,$$

$$H(\boldsymbol{\beta}, \boldsymbol{\alpha}_g) = \frac{\partial S(\boldsymbol{\beta})}{\partial \boldsymbol{\alpha}_g'} = -\sum_{i=1}^{m} \tau_{ig} \sum_{j \in s_i} \left\{ \frac{\partial h[\eta_{ijg}]}{\partial \eta_{ijg}} \right\} \boldsymbol{x}_{ij} \boldsymbol{w}_{ij}', \quad g = 1, \ldots, G,$$

$$H(\boldsymbol{\beta}, \boldsymbol{\beta}) = \frac{\partial S(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}'} = -\sum_{i=1}^{m} \sum_{g=1}^{G} \tau_{ig} \sum_{j \in s_i} \left\{ \frac{\partial h(\eta_{ijg})}{\partial \eta_{ijg}} \right\} \boldsymbol{x}_{ij} \boldsymbol{x}_{ij}',$$

$$H(\pi_g, \pi_l) = \frac{\partial S(\pi_g)}{\partial \pi_g} = -\sum_{i=1}^{m} \left[\frac{\tau_{ig}}{\pi_g^2} \mathbb{1}(g = l) + \frac{\tau_{iG}}{\pi_G^2} \right], \quad g = 1, \ldots, G-1.$$

The remaining (non-redundant) elements of the expected complete data Hessian matrix, that is $H(\boldsymbol{\alpha}_g, \boldsymbol{\alpha}_k)$, $H(\boldsymbol{\beta}, \pi_g)$, and $H(\boldsymbol{\alpha}_g, \pi_g)$, are all null due to parameter distinctiveness.

B Model derivatives for Bernoulli responses

In this section, we provide explicit formulas for computing model derivatives in the case of binary data. As before, let $\tilde{p}_{i(h)}$ denote the np-BP of p_i, conditional on $y_{i\cdot} = h$, with $y_{i\cdot} = \sum_{j \in s_i} y_{ij}$ and let $\tau_{ig(h)}$ be the posterior probability for the i-th small area to belong to the g-th component of the finite mixture, conditional on $y_{i\cdot} = h$. This latter quantity is defined as

$$\tau_{ig(h)} = \frac{\exp\left[\alpha_g h - \sum_{j \in s_i} \log\left(1 + e^{\eta_{ijg}}\right)\right] \pi_g}{\sum_{l=1}^{G} \exp\left[\alpha_l h - \sum_{j \in s_i} \log\left(1 + e^{\eta_{ijl}}\right)\right] \pi_l}.$$

The first derivative vector of $\tilde{p}_{i(h)}$ with respect to model parameters has elements:

$$\begin{aligned}
\frac{\partial \tilde{p}_{i(h)}^{\text{np-BP}}}{\partial \alpha_g} &= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \frac{\partial}{\partial \alpha_g} \left(\sum_{l=1}^{G} p_{ijl}\, \tau_{il(h)} \right) \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \left[\tau_{ig(h)} \frac{\partial p_{ijg}}{\partial \alpha_g} \right] + \left[\sum_{l=1}^{G} p_{ijl} \frac{\partial \tau_{il(h)}}{\partial \alpha_g} \right] \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \{ \tau_{ig(h)} p_{ijg}(1 - p_{ijg}) + \\
&+ \sum_{l=1}^{G} p_{ijl} \left[\frac{\pi_l f_{il(h)} S_{il(h)}(\alpha_l)}{\sum_{k=1}^{G} \pi_k f_{ik(h)}} \mathbb{1}(l = g) - \frac{\pi_l f_{il(h)} \pi_g f_{ig(h)} S_{ig(h)}(\alpha_g)}{\left[\sum_{k=1}^{G} \pi_k f_{ik(h)}\right]^2} \right] \Bigg\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \tau_{ig(h)} p_{ijg}(1 - p_{ijg}) + \sum_{l=1}^{G} p_{ijl}\, \tau_{il(h)} S_{il(h)}(\alpha_l) \mathbb{1}(l = g) - \sum_{l=1}^{G} p_{ijl} \tau_{il(h)} \tau_{ig(h)} S_{ig(h)}(\alpha_g), \right\}.
\end{aligned}$$

with $f_{ig(h)}$ and $S_{ig(h)}(\alpha_g) = h - \sum_{r \in s_i} p_{irg}$ denoting the joint conditional distribution and the score function for the i-th small area, respectively, both conditional on the g-th component of the finite mixture and $y_{i\cdot} = h$.

For the fixed parameters $\boldsymbol{\beta}$, model derivatives are obtained as follows:

$$
\begin{aligned}
\frac{\partial \tilde{p}^{\text{np-BP}}_{i(h)}}{\partial \boldsymbol{\beta}} &= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \frac{\partial}{\partial \boldsymbol{\beta}} \left(\sum_{l=1}^{G} p_{ijl} \tau_{il(h)} \right) \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \left[\sum_{l=1}^{G} \tau_{il(h)} \frac{\partial p_{ijl}}{\partial \boldsymbol{\beta}} \right] + \left[\sum_{l=1}^{G} p_{ijl} \frac{\partial \tau_{il(h)}}{\partial \boldsymbol{\beta}} \right] \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \left[\sum_{l=1}^{G} \tau_{il(h)} p_{ijl} (1 - p_{ijl}) \boldsymbol{x}_{ij} \right] + \right. \\
&\left. + \sum_{l=1}^{G} p_{ijl} \left[\frac{\pi_l f_{il(h)} S_{il(h)}(\boldsymbol{\beta})}{\sum_{k=1}^{G} \pi_k f_{ik(h)}} - \frac{\pi_l f_{il(h)} \sum_{u=1}^{G} \pi_u f_{iu(h)} S_{iu(h)}(\boldsymbol{\beta})}{\left[\sum_{k=1}^{G} \pi_k f_{ik(h)} \right]^2} \right] \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \sum_{l=1}^{G} \tau_{il(h)} p_{ijl} (1 - p_{ijl}) \boldsymbol{x}'_{ij} + \sum_{l=1}^{G} p_{ijl} \tau_{il(h)} S_{il(h)}(\boldsymbol{\beta}) + \right. \\
&\left. - \sum_{l=1}^{G} p_{ijl} \tau_{il(h)} \sum_{k=1}^{G} \tau_{ik(h)} S_{ik(h)}(\boldsymbol{\beta}) \right\},
\end{aligned}
$$

with $S_{ig(h)}(\boldsymbol{\beta}) = -\sum_{r \in s_i} p_{irg} \boldsymbol{x}_{ir}$. Last, for the component probabilities π_g, we have:

$$
\begin{aligned}
\frac{\partial \tilde{p}^{\text{np-BP}}_{i(h)}}{\partial \pi_g} &= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \sum_{l=1}^{G} p_{ijl} \left[\frac{f_{il(h)}}{\sum_{k=1}^{G} \pi_k f_{ik(h)}} \mathbb{1}(g = l) - \frac{f_{il(h)} \pi_l f_{ig(h)}}{(\sum_{k=1}^{G} \pi_k f_{ik(h)})^2} \right] \right\} \\
&= \frac{1}{N_i} \sum_{j=1}^{N_i} \left\{ \sum_{l=1}^{G} p_{ijl} \left[\frac{\tau_{il(h)}}{\pi_l} \mathbb{1}(g = l) - \frac{\tau_{il(h)} \tau_{ig(h)}}{\pi_g} \right] \right\}.
\end{aligned}
$$